
Scalable Bayesian Modelling of Paired Symbols

Ulrich Paquet
Microsoft Research
Cambridge, United Kingdom

Noam Koenigstein
Microsoft R&D
Herzliya, Israel

Ole Winther
Technical University of Denmark
Lyngby, Denmark

Abstract

We present a novel, scalable and Bayesian approach to modelling the occurrence of pairs of symbols (i, j) drawn from a large vocabulary. Observed pairs are assumed to be generated by a simple popularity based selection process followed by censoring using a preference function. By basing inference on the well-founded principle of variational bounding, and using new site-independent bounds, we show how a scalable inference procedure can be obtained for large data sets. State of the art results are presented on real-world movie viewing data.

1 Introduction

We wish to model the occurrence of pairs of discrete symbols (i, j) from a finite set, or predict the occurrence of symbol j given that the other symbol is i . These pairs might be tuples of $(user, item)$ purchase events, or a stream of $(user, game)$ gameplay events. From such a model, a recommender system can be tailored around the conditional probability of item or game j , given user i . Alternatively, these tuples might be $(word_1, word_2)$ bigrams in a simple language model. If there are I and J of each symbol, their discrete density can be fully modelled as a multinomial distribution with $I \times J$ normalized counts, one for each pair. In practice, data is typically sparse compared to large symbol vocabulary sizes, with $I \approx 10^7$ and $J \approx 10^5$ in tasks considered in this paper, and this prevents the full multinomial from generalizing: from user i watching only one movie j , we would like to infer the odds of her viewing other movies j' . This necessitates more compactly parameterized models, which commonly associate real-valued vectors $\mathbf{u}_i \in \mathbb{R}^K$ and $\mathbf{v}_j \in \mathbb{R}^K$ (where $K \ll I, J$) with user i and item j , and draws on an energy $\mathbf{u}_i^T \mathbf{v}_j$ to couple them.

This paper proposes a new approach to modelling the occurrence of pairs of symbols, and makes two main contributions. First, pairwise data is modelled through a simple selection process followed by a preference function that censors the data: in the generative process, pairs (i, j) are presented to a censor at a basic rate, which then chooses to include them in the stream of data with odds that depend only on the coupling energy $\mathbf{u}_i^T \mathbf{v}_j$. Inference is based on the well-founded principle of variational bounding. Second, we show how a scalable procedure can be obtained by using novel looser site-independent bounds.

To see why scalability is a challenge, consider the bilinear softmax distribution

$$p(i, j) = e^{\mathbf{u}_i^T \mathbf{v}_j} / \sum_{i', j'} e^{\mathbf{u}_{i'}^T \mathbf{v}_{j'}} , \quad (1)$$

whose normalizing constant sums over all $I \times J$ discrete options. When i is given, $p(j|i)$ defines softmax regression, the multi-class extension of logistic regression. The bilinear softmax function poses a practical difficulty: the large sums from the normalizing constant appear in the likelihood gradient through $\partial \log p(i, j) / \partial \mathbf{u}_i = \mathbf{v}_j - \sum_{j'=1}^J w_{ij'} \mathbf{v}_{j'}$, where $w_{ij} \doteq e^{\mathbf{u}_i^T \mathbf{v}_j} / \sum_{i', j'} e^{\mathbf{u}_{i'}^T \mathbf{v}_{j'}}$ requires a sum over all IJ pairs in its normalizer. On observing a pair (i, j) , the likelihood is increased by pulling \mathbf{u}_i towards \mathbf{v}_j , while simultaneously pushing it further from all other $\mathbf{v}_{j'}$. There were recent approaches to using the softmax function at scale. Mnih and Teh [15] used noise contrastive estimation [6] to estimate the expensive softmax gradients when training neural probabilistic language

models, which improves on using importance sampling for gradient estimation [2]. In a different approach, the normalization problem can be addressed by redefining $p(j|i)$ as a tree-based hierarchy of smaller softmax functions; this has a direct application to implicit-feedback collaborative filtering [16]. Alternatively, modelling can be done by formulating a simpler objective function based on a classification likelihood, and including stochastically “negative sampled” pairs during optimization. This was done for skip-gram models that consider $(word_1, word_2)$ pairs [13], and for $(user, item)$ pairs [18], where the latter work assumed that each pair can appear at most once. There additionally exists a large body of tangential work, which models an i.i.d. observation *given* i and j , instead of doing density estimation as described above. These include the stochastic block model and its extensions for binary matrices or graphs [1] and the family of “probabilistic matrix factorization” models for a variety of likelihood functions for the observation [5, 12, 14]. The restriction of each pair to appearing at most once places us in the domain of one-class matrix completion [7, 17, 23], where modelling is typically done by formulating different loss functions over the absent pairs (or missing values in the matrix). In these set-ups, a cost value is typically associated with each possible pair. These can be predefined [7, 17] or optimized for [23].

This paper has large-scale collaborative filtering and recommender systems in mind, and places two requirements on the model and inference procedure that do not coexist in other work. **1.** Crucially, *inference should scale with D* , the size of the dataset, i.e. the number of observed pairs, and not with IJ , the number of possible pairs. **2.** We prefer a Bayesian approach that incorporates parameter uncertainty in our inference. This is particularly useful when data is scarce; if game j was played by a handful of users, its lack of usage should be reflected in the posterior estimate of parameters associated with j . To fulfil these requirements, we borrow an unorthodox idea from [18, 19], which views the stream of data as a censored one. Their perspective is that of a one-class model, which contrasts the observations against an unobserved “*negative background*”, although unlike [18], a pair (i, j) can repeatedly be observed. In Section 4, this “negative background” will be employed in various caches as part of the inference pipeline. Our approach practically improves on that of [18], where the data set size was effectively doubled as the non-revealed stream was stochastically resampled and averaged over. As the “non-revealed half” continually changed due to resampling, the inference procedure also did not comfortably map to a distributed architecture. Because exact inference in our model is not possible, we resort to approximating the parameter posterior via a variational lower bound. In this Variational Bayes setting, with a fully factorized posterior approximation, the bound is iteratively maximized through closed-form updates of each factor. The updates are in terms of natural gradients, and are embarrassingly parallel. Empirically, our approach achieves state of the art results on two large-scale recommendation problems (Section 5).

2 Generative model for pairs with censoring

A pair (i, j) will be represented as a pair (\mathbf{y}, \mathbf{z}) of binary indicator vectors, where only bits i and j are “on” in $\mathbf{y} \in \{0, 1\}^I$ and $\mathbf{z} \in \{0, 1\}^J$ respectively. We shall model the data stream by appending a binary variable $o = \text{T}$ (true) to each pair: we *did* observe that symbols i and j co-occurred, user i played game j today, and so on. We therefore observe a stream of D pairs, which takes the form $\{o_d = \text{T}, \mathbf{y}_d, \mathbf{z}_d\}_{d=1}^D$. The censored approach assumes that there were a number of pairs that did not surface in the data stream, such that $o = \text{F}$ (false). *We do not know which pairs and how many they were*, but in practice we will allow the length of the censored stream be specified as a hyperparameter D' , and assume that $\{o_{d'} = \text{F}\}_{d'=1}^{D'}$ is additionally provided. Let data $\mathcal{D} \doteq \{\{o_d = \text{T}, \mathbf{y}_d, \mathbf{z}_d\}_{d=1}^D, \{o_{d'} = \text{F}\}_{d'=1}^{D'}\}$ denote all observations. The ratio D/D' can be seen as a pre-specified positive to negative class ratio; various settings of r in $D' = rD$ are investigated in Section 5. The censored stream constitutes the “negative background” against which the energy $\mathbf{u}_i^T \mathbf{v}_j$ will be fit, and it plays a role similar to that of the softmax normalizer in the gradient of $\log p(i, j)$ from (1): on observing a pair (i, j) , \mathbf{u}_i is pulled towards \mathbf{v}_j and pushed further from all other $\mathbf{v}_{j'}$.

We additionally associate real-valued biases b_i (and b_j) with each user and item, modifying the energy to $\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j$. They play a useful interpretive role in distinguishing between polarizing and non-polarizing content in a recommender system: content that appeals to a wide range of tastes is described by a \mathbf{v}_j with smaller norm, and their appeal is modelled by a positive taste-independent bias. Polarizing content is described by a large-normed \mathbf{v}_j and a negative taste-independent bias; it is only enjoyed by a narrow sliver of tastes.

We propose a model which combines popularity-based selection with a personalized preference function to model (i, j) . **1.** In a *selection step* a user i is chosen with probability π_i , and an item j is chosen with probability ψ_j . **2.** In a *censoring step* the pair (i, j) is observed with probability $\sigma(\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j)$ and censored with probability $1 - \sigma(\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j)$, where $\sigma(a) = 1/(1 + e^{-a})$ is the logistic function.




Figure 1: A generative model for observing D pairs of symbols, assuming that D' unknown pairs were censored.

Let $\mathbf{U} \doteq \{\mathbf{u}_i\}_{i=1}^I$ and $\mathbf{V} \doteq \{\mathbf{v}_j\}_{j=1}^J$ denote all bilinear parameters and $\mathbf{b} \doteq \{\{b_i\}_{i=1}^I, \{b_j\}_{j=1}^J\}$ denote biases, with $\zeta \doteq \{\mathbf{U}, \mathbf{V}, \mathbf{b}\}$. Lastly $\vartheta \doteq \{\zeta, \pi, \psi\}$ includes multinomial parameters π and ψ . The generative process is illustrated in Figure 1, and is as follows: draw parameters ϑ from their prior distributions (given explicitly below). Repeat drawing pairs (i, j) with indexes drawn from $\text{Discrete}(\pi)$ and $\text{Discrete}(\psi)$ and observe the pairs with probability $\sigma(\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j)$. D such pairs are seen, while we assume that D' , the number of censored data points, is specified as a hyperparameter. The density of an uncensored data point d is therefore

$$p(o_d = T, \mathbf{y}_d, \mathbf{z}_d | \vartheta) = p(o_d = T | \mathbf{y}_d, \mathbf{z}_d, \zeta) p(\mathbf{y}_d | \pi) p(\mathbf{z}_d | \psi) = \prod_{i,j} [\pi_i \psi_j \sigma(\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j)]^{y_{di} z_{dj}},$$

while $p(o_{d'} = F | \mathbf{y}_{d'}, \mathbf{z}_{d'}, \zeta) = \prod_{i,j} (1 - \sigma(\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j))^{y_{d'i} z_{d'j}}$ is the odds of censoring pair d' if its indexes were known. The censored indexes $\mathbf{y}_{d'}$ and $\mathbf{z}_{d'}$ are unknown; by including their prior and marginalizing over them, $p(o_{d'} = F | \vartheta)$ is a mixture of IJ components.

The joint density of \mathcal{D} and the unobserved variables $\theta \doteq \{\vartheta, \{\mathbf{y}_{d'}, \mathbf{z}_{d'}\}_{d'=1}^{D'}\}$ depends on further priors on ϑ , for which we choose Dirichlet priors for $p(\pi) = \mathcal{D}(\pi; \alpha_0)$ and $p(\psi) = \mathcal{D}(\psi; \alpha_0)$. The other priors are fully factorized Gaussians, with $p(\mathbf{U}) = \prod_i \mathcal{N}(\mathbf{u}_i; \mathbf{0}, \tau_u^{-1} \mathbf{I})$ and $p(\mathbf{V}) = \prod_j \mathcal{N}(\mathbf{v}_j; \mathbf{0}, \tau_v^{-1} \mathbf{I})$ and, with some overloaded notation, $p(\mathbf{b}) = \prod_i \mathcal{N}(b_i; 0, \tau_b^{-1}) \prod_j \mathcal{N}(b_j; 0, \tau_b^{-1})$. The hierarchical model could be extended further with Gamma hyperpriors on the various Gaussian precisions τ , or Normal-Wishart hyperpriors on both of the Gaussian parameters [20, 22]. If the symbols i and j were accompanied by meta-data tags, these could also be incorporated into the Bayesian model [9]. For the sake of clarity, we omit these additions in this paper. The joint density decomposes as

$$\begin{aligned} p(\mathcal{D}, \theta) &= p(\mathcal{D} | \{\mathbf{y}_{d'}, \mathbf{z}_{d'}\}, \vartheta) p(\{\mathbf{y}_{d'}, \mathbf{z}_{d'}\} | \pi, \psi) p(\vartheta) \\ &= \prod_{i,j} \sigma(\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j)^{c_{ij}} [1 - \sigma(\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j)]^{\sum_{d'} y_{d'i} z_{d'j}} \\ &\quad \cdot \prod_i \pi_i^{c_i + \sum_{d'} y_{d'i}} \cdot \prod_j \psi_j^{c_j + \sum_{d'} z_{d'j}} \cdot p(\mathbf{U}) p(\mathbf{V}) p(\mathbf{b}) p(\pi) p(\psi), \end{aligned} \quad (2)$$

where the uncensored data likelihood was regrouped using observation counts $c_{ij} \doteq \sum_d y_{di} z_{dj} \in \{0, 1, 2, \dots, D\}$ for each pair (i, j) , and marginal counts $c_i \doteq \sum_d y_{di}$ and $c_j \doteq \sum_d z_{dj}$. Note that

$\sum_{i,j} c_{ij} = D$. Marginalizing $p(\mathcal{D}, \theta)$ over $\{\mathbf{y}_{d'}, \mathbf{z}_{d'}\}$ gives a mixture of $\binom{D'+IJ-1}{IJ-1}$ components, each representing a different way of assigning D' indistinguishable F's to IJ distinguishable bins, or assigning nonnegative counts c'_{ij} with $\sum_{i,j} c'_{ij} = D'$ to a “negative class count matrix”.

At first glance of (2), it would seem as if inference would still scale with IJ , and that we have done nothing more than match the bilinear softmax's $\mathcal{O}(IJ)$ computational burden through the introduction of D' . The following sections are devoted to developing a variational approximation, and with it a practically scalable inference scheme that relies on various “negative background” caches.

3 Variational Bayes

To find a scalable yet Bayesian inference procedure, we approximate $p(\theta|\mathcal{D})$ with a factorized surrogate density $q(\theta)$, found by maximizing a variational lower bound to $\log p(\mathcal{D})$ [24]. First, we lower-bound each logistic function in (2) by associating a parameter ξ_{ij} with it [8]. Dropping subscripts, each bound would be $\sigma(\pm a) \geq \sigma(\xi) \exp(-\lambda(\xi) (a^2 - \xi^2) \pm \frac{a}{2} - \frac{\xi}{2})$, where the lower bound on $1 - \sigma(a)$ is that of $\sigma(-a)$ above. The bound depends on the deterministic function $\lambda(\xi) \doteq \frac{1}{2\xi} [\sigma(\xi) - \frac{1}{2}]$. Let $\xi \doteq \{\xi_{ij}\}$ denote the set of logistic variational parameters, and substitute the bound into (2) to get $p(\mathcal{D}, \theta) \geq p_\xi(\mathcal{D}, \theta)$. Our variational objective $\mathcal{L}_\xi[q]$, as a function of ξ and functional of q , follows from

$$\log p(\mathcal{D}) \geq \log \int p_\xi(\mathcal{D}, \theta) d\theta \geq \int q(\theta) \log \frac{p_\xi(\mathcal{D}, \theta)}{q(\theta)} d\theta \doteq \mathcal{L}_\xi[q], \quad (3)$$

which will be maximized with respect to q and ξ . The factorization of q employed in this paper is

$$q(\theta) \doteq \prod_i q(b_i) \prod_k q(u_{ik}) \cdot \prod_j q(b_j) \prod_k q(v_{jk}) \cdot \prod_{d'} q(\mathbf{y}_{d'}) q(\mathbf{z}_{d'}) \cdot q(\pi) q(\psi). \quad (4)$$

The factors approximating each symbol's features in \mathbf{U} , \mathbf{V} , and \mathbf{b} are chosen to be a Gaussian, for example $q(u_{ik}) = \mathcal{N}(u_{ik}; \mu_{ik}, \omega_{ik}^{-1})$. The approximating factors $q(\pi)$ and $q(\psi)$ are both Dirichlet, for example $q(\pi) = \mathcal{D}(\pi; \alpha)$. The bound in (3) is stated fully in Appendix B.

For the purpose of obtaining a scalable algorithm, the most important parameterizations are for the categorical (discrete) factors $q(\mathbf{y}_{d'})$ and $q(\mathbf{z}_{d'})$. We shall argue and show in Sections 4 and 5 that choosing $D' \approx D$ is desired, and as D' is potentially large, the *parameters of $q(\mathbf{y}_{d'})$ will be tied*. *This tying of parameters is the key to achieving a scalable algorithm*. We let all $q(\mathbf{y}_{d'})$ share the same parameter vector \mathbf{s} on the probability simplex, such that $q(\mathbf{y}_{d'}) = \prod_i s_i^{y_{d'i}}$ for all d' . Similarly, all $q(\mathbf{z}_{d'})$ share probability vector \mathbf{t} , such that $q(\mathbf{z}_{d'}) = \prod_j t_j^{z_{d'j}}$ for all d' .

Making and trading predictions Our original desideratum was to infer the probability of symbol j , conditional on the other symbol being i , and the observed data. Bayesian marginalization requires us to average the predictions over the model parameter posterior distribution. Here it is an analytically intractable task, which we approximate by using q as a surrogate for the true posterior. Firstly, $p(o = \mathsf{T} | \mathbf{y}, \mathbf{z}, \mathcal{D}) \approx \int p(o = \mathsf{T} | \mathbf{y}, \mathbf{z}, \theta) q(\theta) d\theta = \int \sigma(a_{ij}) \mathcal{N}(a_{ij}; \mu_{ij}, \sigma_{ij}^2) da_{ij} \approx \sigma(x_{ij})$ if $y_i = z_j = 1$. The random variable a_{ij} was defined as $a_{ij} \doteq \mathbf{u}_i^T \mathbf{v}_j + b_i + b_j$, with its density approximated with its first two moments under q , i.e. $\mu_{ij} \doteq \mathbb{E}_q[a_{ij}]$ and $\sigma_{ij}^2 \doteq \mathbb{E}_q[(a_{ij} - \mu_{ij})^2]$. With $x_{ij} \doteq \mu_{ij} / (1 + \pi \sigma_{ij}^2 / 8)^{1/2}$, the final approximation of a logistic Gaussian integral follows from [10]. Again using q , the posterior density of symbol j , provided that the first symbol is i , is approximately proportional to (writing “ T ” for “ $o = \mathsf{T}$ ” for brevity)

$$p(z_j = 1 | \mathsf{T}, y_i = 1, \mathcal{D}) \propto p(\mathsf{T} | y_i = z_j = 1, \mathcal{D}) \int p(z_j = 1 | \psi) q(\psi) d\psi = \sigma(x_{ij}) \mathbb{E}_q[\psi_j]. \quad (5)$$

Hence $p(z_j = 1 | o = \mathsf{T}, y_i = 1, \mathcal{D}) \approx \sigma(x_{ij}) \mathbb{E}_q[\psi_j] / \sum_{j'} \mathbb{E}_q[\psi_{j'}] \sigma(x_{ij'})$, normalizing to one.

4 Scalable inference

A scalable update procedure for the factors of $q(\theta)$ is presented in this section, culminating in Algorithm 1. The algorithm optimizes over t_{\max} loops, but can also be run until complete convergence as

Algorithm 1: Paired Symbol Modelling

input: \mathcal{D} (or D' and all non-zero c_{ij}), $\alpha_0, \beta_0, \tau_u, \tau_v, \tau_b$

initialize: $\xi^* \leftarrow 1, \mathbf{s} \leftarrow [1/I], \mathbf{t} \leftarrow [1/J]$

for $t = 1 : t_{\max}$ **do**

update $q(\pi)$; update $q(\psi)$; cache item-background $\mathbf{P}_\ominus, \mathbf{m}_\ominus^\dagger, \mathbf{m}_\ominus^\ddagger, \nu_\ominus, \varkappa_\ominus$; update \mathbf{s}

pfor $i = 1 : I$ **do** { update $q(b_i)$; update $\prod_{k=1}^K q(u_{ik})$ }

cache user-background $\mathbf{P}_\oplus, \mathbf{m}_\oplus^\dagger, \mathbf{m}_\oplus^\ddagger, \nu_\oplus, \varkappa_\oplus$; update ξ^* ; update \mathbf{t}

pfor $j = 1 : J$ **do** { update $q(b_j)$; update $\prod_{k=1}^K q(v_{jk})$ }

the evidence lower bound \mathcal{L} from (3) can be explicitly calculated. We use **pfor** to indicate embarrassingly parallel loops, although the updates for \mathbf{s} , \mathbf{t} , and ξ^* also make extensive use of parallelization.

Let graph $\mathcal{G} = \{(i, j) : c_{ij} > 0\}$ be the sparse set of all observed pair indexes. As there are IJ logistic variational parameters ξ_{ij} , we shall divide them into two sets, those with indexes in \mathcal{G} , and those without. Therefore ξ_{ij} shall be optimized for when $(i, j) \in \mathcal{G}$, while the ξ_{ij} 's shall *share the same parameter value* ξ^* for $(i, j) \notin \mathcal{G}$. Even though the form of (2) suggests that we would need two versions of ξ_{ij} , one for the bounded σ -term, and one its opposite, this is not required, as the optimization of the bound is symmetric. When ξ_{ij} maximizes \mathcal{L} on the bounded σ -term, it simultaneously maximizes \mathcal{L} on the bounded $(1 - \sigma)$ -term. We'll use the shorthand $\lambda_{ij} \doteq \lambda(\xi_{ij})$ for $(i, j) \in \mathcal{G}$; similarly, λ^* denotes $\lambda(\xi^*)$ when $(i, j) \notin \mathcal{G}$. The updates for symbols i and j 's parameters mirror each other, and only the ‘‘user updates’’ are laid out in this section.

Gaussian updates for $q(u_{ik})$ We will present here a bulk update of $\prod_k q(u_{ik})$, which is faster than sequentially maximizing $\mathcal{L}_\xi[q]$ with respect to each of them in turn. We first solve for the maximum of \mathcal{L} with respect to a full Gaussian (not factorized) approximation $\tilde{q}(\mathbf{u}_i) = \mathcal{N}(\mathbf{u}_i; \boldsymbol{\mu}_i, \mathbf{P}_i^{-1})$. The fully factorized $q(u_{ik})$ can then be recovered from the intermediate approximation $\tilde{q}(\mathbf{u}_i)$ as those that minimize the Kullback-Leibler divergence $D_{\text{KL}}(\prod_k q(u_{ik}) \parallel \tilde{q}(\mathbf{u}_i))$: this is achieved when the means of $q(u_{ik})$ match that of $\tilde{q}(\mathbf{u}_i)$, while their precisions match the diagonal precision of $\tilde{q}(\mathbf{u}_i)$. The validity of the intermediate bound is proved in Appendix B.2. The updates rely on careful caching, which we'll first illustrate through \tilde{q} 's precision matrix. \mathcal{L} is maximized when $\tilde{q}(\mathbf{u}_i)$ has as natural parameters a precision matrix

$$\mathbf{P}_i = \sum_{j \in \mathcal{G}(i)} c_{ij} \cdot 2\lambda_{ij} \cdot \mathbb{E}_q[\mathbf{v}_j \mathbf{v}_j^T] + \overbrace{\sum_{d'} \sum_j \mathbb{E}_q[y_{d'i} z_{d'j}] \cdot 2\lambda_{ij} \cdot \mathbb{E}_q[\mathbf{v}_j \mathbf{v}_j^T]}^{(a)} + \tau_u \mathbf{I} \quad (6)$$

and a mean-times-precision vector \mathbf{m}_i , which we will state later. Looking at \mathbf{P}_i in (6), an undesirable sum over all d' and j is required in (a). We endeavoured that the update would be *sparse*, and only sum over observed indexes in $\mathcal{G}(i) \doteq \{j : (i, j) \in \mathcal{G}\}$. The benefit of the shared variational parameters now becomes apparent. With $\mathbb{E}_q[y_{d'i} z_{d'j}] = s_i t_j$ and $\lambda_{ij} = \lambda^*$ when $(i, j) \notin \mathcal{G}$, the sum in (a) decomposes as

$$(a) = \sum_{j \in \mathcal{G}(i)} s_i t_j D' \cdot 2(\lambda_{ij} - \lambda^*) \mathbb{E}_q[\mathbf{v}_j \mathbf{v}_j^T] + s_i D' \cdot 2\lambda^* \cdot \overbrace{\sum_j t_j \mathbb{E}_q[\mathbf{v}_j \mathbf{v}_j^T]}^{\text{negative background } \mathbf{P}_\ominus}.$$

Barring the ‘‘negative background’’ term, only a sparse sum that involves observed pairs is required. This background term is rolled up into a global *item*-background cache, which is computed once before updating all $q(u_{ik})$. Throughout the paper, the \ominus symbol will denote an item-background cache. The cache $\mathbf{P}_\ominus \doteq \sum_j t_j \mathbb{E}_q[\mathbf{v}_j \mathbf{v}_j^T]$ is used in each precision matrix update, for example

$$\mathbf{P}_i = s_i D' \cdot 2\lambda^* \cdot \mathbf{P}_\ominus + \sum_{j \in \mathcal{G}(i)} \left(c_{ij} \cdot 2\lambda_{ij} + s_i t_j D' \cdot 2(\lambda_{ij} - \lambda^*) \right) \mathbb{E}_q[\mathbf{v}_j \mathbf{v}_j^T] + \tau_u \mathbf{I}.$$

We've deliberately laboured the above decomposition of an expensive update into a background cache and a sparse sum over actual observations, as it serves as a template for other parameter

updates to come. Turning to the mean-times-precision vector $\mathbf{m}_i \doteq \mathbf{P}_i \boldsymbol{\mu}_i$ of $\tilde{q}(\mathbf{u}_i)$, we find that

$$\mathbf{m}_i = \mathbb{E}_q \left[\sum_{j \in \mathcal{G}(i)} c_{ij} \left(\frac{1}{2} - 2\lambda_{ij}(b_i + b_j) \right) \mathbf{v}_j + \sum_{d'} \sum_j y_{d'i} z_{d'j} \left(-\frac{1}{2} - 2\lambda_{ij}(b_i + b_j) \right) \mathbf{v}_j \right]. \quad (7)$$

There is a subtle link between (7) and the gradients of the bilinear soft-max likelihood, which we'll explore in the next paragraph. To find \mathbf{m}_i , two additional caches are added to the item-background cache, and are computed once before any $q(u_{ik})$ updates. They are $\mathbf{m}_\ominus^\dagger \doteq \sum_j t_j \mathbb{E}_q[b_j] \mathbb{E}_q[\mathbf{v}_j]$ and $\mathbf{m}_\ominus^\ddagger \doteq \sum_j t_j \mathbb{E}_q[\mathbf{v}_j]$. The final mean-times-precision update is

$$\begin{aligned} \mathbf{m}_i &= s_i D' \left[\left(-\frac{1}{2} - 2\lambda^* \mathbb{E}_q[b_i] \right) \mathbf{m}_\ominus^\dagger - 2\lambda^* \mathbf{m}_\ominus^\ddagger \right] \\ &\quad + \sum_{j \in \mathcal{G}(i)} \left(c_{ij} \left(\frac{1}{2} - 2\lambda_{ij} \mathbb{E}_q[b_i + b_j] \right) - s_i t_j D' \cdot 2(\lambda_{ij} - \lambda^*) \mathbb{E}_q[b_i + b_j] \right) \mathbb{E}_q[\mathbf{v}_j], \end{aligned} \quad (8)$$

and again only sums over $j \in \mathcal{G}(i)$ and not all J . There are of course additional variational parameters ξ_{ij} , and they are computed and discarded when needed according to (11).

Bilinear softmax gradients The connection between this model and a bilinear softmax model can be seen when the biases are ignored. Consider the gradient of \mathcal{L} with respect to *mean* parameter $\boldsymbol{\mu}_i$,

$$\nabla \mathcal{L}(\boldsymbol{\mu}_i) = -\mathbf{P}_i \boldsymbol{\mu}_i + \frac{1}{2} \left(\sum_{j \in \mathcal{G}(i)} c_{ij} \mathbb{E}_q[\mathbf{v}_j] - D' \sum_j s_i t_j \mathbb{E}_q[\mathbf{v}_j] \right). \quad (9)$$

The gradient $\nabla \mathcal{L}(\boldsymbol{\mu}_i)$ is zero at (7), which was stated, together with (6), in terms of *natural* parameters. As $\mathcal{L}(\boldsymbol{\mu}_i)$ is quadratic, it can be exactly maximized; furthermore, the maximum with respect to \mathbf{P}_i is attained at the negative Hessian $\mathbf{P}_i = -\nabla^2 \mathcal{L}(\boldsymbol{\mu}_i)$, given in (6). The curvature of the bound, as a function of $\boldsymbol{\mu}_i$, directly translates into our posterior approximation's uncertainty of \mathbf{u}_i . The log likelihood of a softmax model would be $L = \sum_d \log p(i_d, j_d)$, with the likelihood of each pair defined by (1). The gradient of the log likelihood is therefore

$$\nabla L(\mathbf{u}_i) = \sum_{j \in \mathcal{G}(i)} c_{ij} \mathbf{v}_j - D \sum_j w_{ij} \mathbf{v}_j, \quad (10)$$

with weights $w_{ij} \doteq e^{\mathbf{u}_i^T \mathbf{v}_j} / \sum_{i', j'} e^{\mathbf{u}_{i'}^T \mathbf{v}_{j'}}$ that sum to one over all IJ options. The weights in (9) were simply $s_i t_j$, and also sum to one over all options. The difference between (9) and (10) is that $s_i t_j$ is used as a *factorized substitute* for w_{ij} . This simplification allows the convenience that none of the updates described in Section 4 need to be stochastic, and substitute functions, as employed by noise contrastive divergence to maximize L , are not required. (The Hessian $\nabla^2 L(\mathbf{u}_i)$ contains a *double-sum* over indexes j .) Considering the two equations above, one might expect to set hyperparameter D' to $D' = D$, and in Section 5 we show that this is a reasonable choice.

Gaussian updates for $q(b_i)$ The maximum of \mathcal{L} with respect to $q(b_i)$ *re-uses* cache $\mathbf{m}_\ominus^\dagger$, but requires the additional cache $\nu_\ominus \doteq \sum_j t_j \mathbb{E}_q[b_j]$ to be precomputed. Gaussian $q(b_i)$ has a mean-times-precision parameter $\nu_i = s_i D' \left(-\frac{1}{2} - 2\lambda^* (\nu_\ominus + \mathbb{E}_q[\mathbf{u}_i^T] \mathbf{m}_\ominus^\dagger) \right) + \sum_{j \in \mathcal{G}(i)} \left(c_{ij} \left(\frac{1}{2} - 2\lambda_{ij} \mathbb{E}_q[\mathbf{u}_i^T \mathbf{v}_j + b_j] \right) - s_i t_j D' \cdot 2(\lambda_{ij} - \lambda^*) \mathbb{E}_q[\mathbf{u}_i^T \mathbf{v}_j + b_j] \right)$, and its precision parameter $\rho_i = 2\lambda^* s_i D' + \sum_{j \in \mathcal{G}(i)} (c_{ij} 2\lambda_{ij} + s_i t_j D' 2(\lambda_{ij} - \lambda^*)) + \tau_b$ follows a similar form.

Logistic bound parameter updates As discussed above, the logistic bound parameters ξ_{ij} associated with observations $(i, j) \in \mathcal{G}$ are treated individually whilst the remainder are shared and denoted by ξ^* . The individually optimized bound

$$\xi_{ij}^2 = \mathbb{E}_q[(\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j)^2] \quad (11)$$

can be used anytime during the updates and then discarded (we always use the positive root for ξ_{ij}). The shared parameter can be written in terms of cached quantities and a sum that scales with D (the user-background cache is denoted with a \oplus symbol, and mirrors the item-background cache):

$$(\xi^*)^2 = \frac{1}{Z} \left(\text{tr } \mathbf{P}_\oplus \mathbf{P}_\ominus + 2\mathbf{m}_\oplus^\dagger \mathbf{m}_\ominus^\dagger + 2\mathbf{m}_\oplus^\dagger \mathbf{m}_\ominus^\ddagger + 2\nu_\oplus \nu_\ominus + \kappa_\oplus + \kappa_\ominus - \sum_{(i,j) \in \mathcal{G}} s_i t_j \xi_{ij}^2 \right)$$

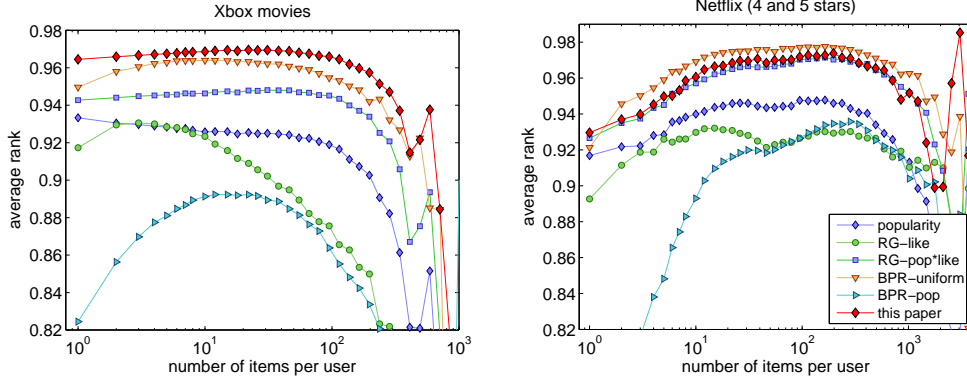


Figure 2: The rank $R(i, j^*)$ in (12), averaged over users and grouped logarithmically by c_i . The *top* evaluation is on the *Xbox movies* sample, while the *bottom* one is on the “implicit feedback” *Netflix (4 and 5 stars)* set.

where $\mathcal{Z} \doteq 1 - \sum_{(i,j) \in \mathcal{G}} s_i t_j$. Cache $\kappa_\Theta \doteq \sum_j t_j \mathbb{E}_q[b_j^2]$ also plays a role in the categorical updates.

Dirichlet updates As the multinomial distribution is conjugate to a Dirichlet, its updates have a particularly simple form. $q(\pi)$ is Dirichlet $\mathcal{D}(\pi; \alpha)$ with parameters $\alpha_i = \alpha_{0i} + c_i + s_i D'$. Each pseudo-count adds c_i , the number of views for user i , to the expected number of views that were censored and not made.

Categorical updates There are D' categorical (discrete) factors $q(\mathbf{y}_{d'})$, and the key to finding a scalable inference procedure lies in tying all their parameters together in \mathbf{s} , with $\sum_i s_i = 1$. Looking at the second line of (2), the factors depend on the expected bounded logistic functions

$$\Omega_{ij} \doteq \log \sigma(\xi_{ij}) - \lambda(\xi_{ij}) \left(\mathbb{E}_q[(\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j)^2] - \xi_{ij}^2 \right) - \frac{\xi_{ij}}{2} - \frac{1}{2} \mathbb{E}_q[\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j].$$

The categorical parameters are, if we solve for all the D' tied distributions $q(\mathbf{y}_{d'})$ jointly,

$$s_i \propto \exp \left(\mathbb{E}_q[\log \pi_i] + \sum_j t_j \Omega_{ij} \right).$$

In practice, each entry $\log s_i + \text{const}$ can be computed in parallel; afterwards, they are renormalized to give \mathbf{s} . To find \mathbf{s} , an efficient way is needed to determine $\sum_j t_j \Omega_{ij}$, and this can again be done with careful bookkeeping. The observed terms $j \in \mathcal{G}(i)$ are treated differently from the rest. For observed terms we can use the optimal logistic parameters in (11) to simplify $\Omega_{ij} \doteq \log \sigma(\xi_{ij}) - \frac{\xi_{ij}}{2} - \frac{1}{2} \mathbb{E}_q[\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j]$. By denoting $\Omega_{ij}(\xi^*)$ evaluated with the shared parameter ξ^* by Ω_{ij}^* , we can write $\sum_j t_j \Omega_{ij} = \sum_{j \in \mathcal{G}(i)} t_j (\Omega_{ij} - \Omega_{ij}^*) + \sum_j t_j \Omega_{ij}^*$. The first term scales with D and the second term can be written using cached quantities: $\sum_j t_j \Omega_{ij}^* = -\lambda^* (\text{tr} \mathbb{E}_q[\mathbf{u}_i \mathbf{u}_i^T] \mathbf{P}_\Theta + 2 \mathbb{E}_q[b_i \mathbf{u}_i^T] \mathbf{m}_\Theta^\dagger + 2 \mathbb{E}_q[\mathbf{u}_i^T] \mathbf{m}_\Theta^\dagger + \mathbb{E}_q[b_i^2] + 2 \mathbb{E}_q[b_i] \nu_\Theta + \kappa_\Theta) + \log \sigma(\xi^*) + \frac{(\xi^*)^2}{2} \lambda^* - \frac{\xi^*}{2} - \frac{1}{2} (\mathbb{E}_q[\mathbf{u}_i^T] \mathbf{m}_\Theta^\dagger + \mathbb{E}_q[b_i] + \nu_\Theta)$.

5 Evaluation

A key application for modelling paired (*user, item*) symbols is large-scale recommendation systems, and we evaluate the predictions obtained by (5) on two large data sets.¹ The *Xbox movies* data is a sample of 5.6×10^7 views for 6.2×10^6 users on a sub-catalogue of around 1.2×10^4 movies [18]. To evaluate on data known in the Machine Learning community, the four- and five-starred ratings from the Netflix prize data set were used to simulate a stream of “implicit feedback” (*user, item*) pairs in the *Netflix (4 and 5 stars)* data. We refer the reader to [18] for a complete data set description. For each user, one item was randomly removed to create a test set. To mimic a real scenario in the simplest possible way, each user’s non-viewed items were ranked, and the position of the test item

¹Additional results follow in the Appendix D.




Figure 3: The average rank $R(i, j^*)$ in (12), grouped logarithmically by c_i , for varying values of r in $D' = rD$.




Figure 4: The empirical densities of $\sigma(x_{j^*})$, as defined in (5), over all held-out items j^* in the *Netflix (4 and 5 stars)* set. The densities are sliced according to $c_i = 1, \dots, 40$ for different values of r in $D' = rD$.

noted. We are interested in the rank of held out item j^* for user i on a $[0, 1]$ scale,

$$R(i, j^*) \doteq \frac{1}{J - |\mathcal{G}(i)|} \sum_{j \notin \mathcal{G}(i)} \mathbb{I}[f_{ij^*} > f_{ij}], \quad (12)$$

where f_{ij} indicates the score given by (5) or any alternative algorithm.

In Figure 2, we facet the average rank by c_i , the number of movie views per user. As the evaluation is over 6 million users, this gives a more representative perspective than reporting a single average. Apart from ranking by *popularity* c_j , which would be akin to only factorizing with $s_i t_j$, we compare against two other baselines. *BPR-uniform* and *BPR-pop* represent different versions of the Bayesian Personalized Ranking algorithm [21], which optimizes a rank metric directly against either the data distribution of items (*BPR-uniform*, with missing items are sampled uniformly during stochastic optimization), or a tilted distribution aimed at personalizing recommendations regardless of an item's popularity (*BPR-pop*, with missing items sampled proportional to their popularity). Their hyperparameters were set using cross-validation. For the Random Graph model [18], rankings are shown with pure personalization (*RG-like*) and with an item popularity rate factored in (*RG-pop*like*). The comparison in Figure 2 is drawn using $K = 20$ dimensions, $D' = D$ and hyperparameters set to one. For *Xbox movies*, the model outperforms all alternatives that we compared against. *BPR-uniform*, optimizing (12) directly, performs slightly better on the less sparse *Netflix* set (the *Xbox* usage sample is much sparser, as it is easier to rate many movies than to view as many). For *Xbox movies*, updating all item-related parameters in Algorithm 1 took 69 seconds on a 24-core (Intel Xeon 2.93Ghz) machine, and updating all user-related parameters took 83 seconds.

We surmised in Section 4 that $D' = D$ is a reasonable hyperparameter setting, and Figure 3 validates this claim. The figure shows the average held-out rank on the *Netflix (4 and 5)* set for various settings of D' through $D' = rD$. The average rank improves beyond $r = 1$, but empirically slowly decreases beyond $r = 2$. To provide insight into the “censoring” step, Figure 4 accompanies Figure 3, and shows the empirical density of the Bernoulli variable $\sigma(x_{j^*})$ for held-out items j^* . We break the empirical density down over users that appear in $c_i = 1, 2, 3, \dots, 40$ pairs. Given that the held-out

pairs were observed, the Bernoulli variable should be *true*, and the density of $\sigma(x_{j*})$ shifts right as c_i becomes bigger. The effect of having to explain less ($r = \frac{1}{2}$) or more ($r = 2$) censored pairs is also visible in the figure. There is also a slight benefit in increasing K . The average rank \hat{R}_{20} for $K = 20$ is 0.9649, using $r = 1$. An increased latent dimensionality gives $\hat{R}_{30} - \hat{R}_{20} = 1.07 \times 10^{-4}$, $\hat{R}_{40} - \hat{R}_{20} = 1.73 \times 10^{-4}$, and $\hat{R}_{50} - \hat{R}_{20} = 0.87 \times 10^{-4}$.

6 Summary and outlook

In this paper we presented a novel model for pairs of symbols, and showed state of the art results on a large-scale movies recommendation task. Scalability was achieved by factorizing the popularity or selection step via $\pi_i \psi_j$, and employing “site-independent” variational bounds through careful parameter tying. This approach might be too simplistic; an extension would be to use a N -component mixture model to select pairs with odds $\sum_{n=1}^N \pi_{in} \psi_{jn}$, and perform inference with Gibbs sampling.

It is worth noting that Böhning [3] and Bouchard [4] provide lower bounds to the logarithm of (1). We originally embarked on a variational approximation to a posterior with (1) as likelihood using Bouchard’s bound, for which bookkeeping like Section 4’s was done. However, with realistically large I and J , solutions were trivial, as the means of the variational posterior approximations for \mathbf{u}_i and \mathbf{v}_j were zero. We leave Böhning’s bound to future work.

A The Joint Model

The joint density in (2) follows from combining the data likelihood

$$\begin{aligned} p(\mathcal{D}|\boldsymbol{\theta}) &= \prod_d p(o_d = \mathbf{T}|\mathbf{y}_d, \mathbf{z}_d, \boldsymbol{\theta}) p(\mathbf{y}_d|\boldsymbol{\pi}) p(\mathbf{z}_d|\boldsymbol{\psi}) \cdot \prod_{d'} p(o_{d'} = \mathbf{F}|\mathbf{y}_{d'}, \mathbf{z}_{d'}, \boldsymbol{\theta}) \\ &= \prod_d \prod_{i,j} [\pi_i \psi_j \sigma(\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j)]^{y_{di} z_{dj}} \prod_{d'} \prod_{i,j} (1 - \sigma(\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j))^{y_{d'i} z_{d'j}} \end{aligned}$$

with a prior on the unobserved variables $\boldsymbol{\theta}$, and rewriting the expression using observation counts $c_{ij} \doteq \sum_d y_{di} z_{dj}$ for each pair (i, j) , and marginal counts $c_i \doteq \sum_d y_{di}$ and $c_j \doteq \sum_d z_{dj}$. The joint density is shown in Figure 1.

B The Variational Bound

For the sake of later derivations, it is worthwhile to explicitly write $\mathcal{L}_{\boldsymbol{\xi}}[q]$ as it appears in (3). It is

$$\begin{aligned} \mathcal{L}_{\boldsymbol{\xi}}[q] &= \sum_{i,j} c_{ij} \mathbb{E}_q \left[\log \sigma(\xi_{ij}) - \lambda(\xi_{ij}) \left((\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j)^2 - \xi_{ij}^2 \right) + \frac{1}{2} (\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j) - \frac{\xi_{ij}}{2} \right] \\ &\quad + \sum_{i,j} \sum_{d'} \mathbb{E}_q[y_{d'i} z_{d'j}] \mathbb{E}_q \left[\log \sigma(\xi_{ij}) - \lambda(\xi_{ij}) \left((\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j)^2 - \xi_{ij}^2 \right) \right. \\ &\quad \quad \quad \left. - \frac{1}{2} (\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j) - \frac{\xi_{ij}}{2} \right] \\ &\quad + \sum_i \left(c_i + \sum_{d'} \mathbb{E}[y_{d'i}] \right) \mathbb{E}_q[\log \pi_i] + \sum_j \left(c_j + \sum_{d'} \mathbb{E}[z_{d'j}] \right) \mathbb{E}_q[\log \psi_j] \\ &\quad + \sum_i \sum_k \mathbb{E}_q[\log p(u_{ik})] + \sum_j \sum_k \mathbb{E}_q[\log p(v_{jk})] + \sum_i \mathbb{E}_q[\log p(b_i)] + \sum_j \mathbb{E}_q[\log p(b_j)] \\ &\quad + \mathbb{E}_q[\log p(\boldsymbol{\pi})] + \mathbb{E}_q[\log p(\boldsymbol{\psi})] \\ &\quad - \sum_i \sum_k \mathbb{E}_q[\log p(u_{ik})] - \sum_j \sum_k \mathbb{E}_q[\log p(v_{jk})] - \sum_i \mathbb{E}_q[\log p(b_i)] - \sum_j \mathbb{E}_q[\log p(b_j)] \end{aligned}$$

$$- \mathbb{E}_q[\log q(\boldsymbol{\pi})] - \mathbb{E}_q[\log q(\boldsymbol{\psi})] - \sum_{d'} \mathbb{E}_q[\log p(\mathbf{y}_{d'})] - \sum_{d'} \mathbb{E}_q[\log p(\mathbf{z}_{d'})]. \quad (13)$$

All expectations are taken under $q(\boldsymbol{\theta})$ defined in (4).

B.1 Bookkeeping

The scalable parameter updates in Section 4 rely on a number of cached quantities, which we state here together for completeness:

$$\left. \begin{aligned} \mathbf{P}_\ominus &\doteq \sum_j t_j \mathbb{E}_q[\mathbf{v}_j \mathbf{v}_j^T] & \mathbf{m}_\ominus^\dagger &\doteq \sum_j t_j \mathbb{E}_q[b_j] \mathbb{E}_q[\mathbf{v}_j] & \mathbf{m}_\ominus^\ddagger &\doteq \sum_j t_j \mathbb{E}_q[\mathbf{v}_j] \\ \nu_\ominus &\doteq \sum_j t_j \mathbb{E}_q[b_j] & \kappa_\ominus &\doteq \sum_j t_j \mathbb{E}_q[b_j^2] \end{aligned} \right\} \text{item-background}$$

$$\left. \begin{aligned} \mathbf{P}_\oplus &\doteq \sum_i s_i \mathbb{E}_q[\mathbf{u}_i \mathbf{u}_i^T] & \mathbf{m}_\oplus^\dagger &\doteq \sum_i s_i \mathbb{E}_q[b_i] \mathbb{E}_q[\mathbf{u}_i] & \mathbf{m}_\oplus^\ddagger &\doteq \sum_i s_i \mathbb{E}_q[\mathbf{u}_i] \\ \nu_\oplus &\doteq \sum_i s_i \mathbb{E}_q[b_i] & \kappa_\oplus &\doteq \sum_i s_i \mathbb{E}_q[b_i^2] \end{aligned} \right\} \text{user-background}$$

B.2 Latent trait vector updates

We stated $q(\boldsymbol{\theta})$ in terms of the factorized Gaussian $\prod_k q(u_{ik})$, and will solve for $\prod_k q(u_{ik})$ by first maximizing an intermediate lower bound with respect to the full-covariance Gaussian $\tilde{q}(\mathbf{u}_i) \doteq \mathcal{N}(\mathbf{u}_i; \boldsymbol{\mu}_i, \mathbf{P}_i^{-1})$. Once $\tilde{q}(\mathbf{u}_i)$ is found, a lower bound to it is maximized to find $\prod_k q(u_{ik})$.

B.2.1 Scalable updates

Let $\lambda_{ij} \doteq \lambda(\xi_{ij})$. The variational bound in (13), as a function of the full-covariance Gaussian $\tilde{q}(\mathbf{u}_i)$, is

$$\begin{aligned} \mathcal{L}[\tilde{q}(\mathbf{u}_i)] &= -\frac{1}{2} \sum_j c_{ij} \left(2\lambda_{ij} \text{tr}(\mathbb{E}_{\tilde{q}}[\mathbf{u}_i \mathbf{u}_i^T] \mathbb{E}_q[\mathbf{v}_j \mathbf{v}_j^T]) - 2\mathbb{E}_{\tilde{q}}[\mathbf{u}_i]^T \left(\frac{1}{2} - 2\lambda_{ij} \mathbb{E}_q[b_i + b_j] \right) \mathbb{E}_q[\mathbf{v}_j] \right) \\ &\quad - \frac{1}{2} \sum_j \sum_{d'} \mathbb{E}_q[y_{d'i}] \mathbb{E}_q[z_{d'j}] \left(2\lambda_{ij} \text{tr}(\mathbb{E}_{\tilde{q}}[\mathbf{u}_i \mathbf{u}_i^T] \mathbb{E}_q[\mathbf{v}_j \mathbf{v}_j^T]) \right. \\ &\quad \left. - 2\mathbb{E}_{\tilde{q}}[\mathbf{u}_i]^T \left(-\frac{1}{2} - 2\lambda_{ij} \mathbb{E}_q[b_i + b_j] \right) \mathbb{E}_q[\mathbf{v}_j] \right) - \frac{1}{2} \text{tr}(\mathbb{E}_q[\mathbf{u}_i \mathbf{u}_i^T] \tau_u \mathbf{I}) - \mathbb{E}_{\tilde{q}}[\log \tilde{q}(\mathbf{u}_i)] \\ &= -\frac{1}{2} \text{tr} \mathbb{E}_{\tilde{q}}[\mathbf{u}_i \mathbf{u}_i^T] \left(\tau_u \mathbf{I} + \sum_j 2\lambda_{ij} \left(c_{ij} + \sum_{d'} \mathbb{E}_q[y_{d'i}] \mathbb{E}_q[z_{d'j}] \right) \mathbb{E}_q[\mathbf{v}_j \mathbf{v}_j^T] \right) \\ &\quad + \mathbb{E}_{\tilde{q}}[\mathbf{u}_i]^T \sum_j \left(c_{ij} \left(\frac{1}{2} - 2\lambda_{ij} \mathbb{E}_q[b_i + b_j] \right) \right. \\ &\quad \left. + \sum_{d'} \mathbb{E}_q[y_{d'i}] \mathbb{E}_q[z_{d'j}] \left(-\frac{1}{2} - 2\lambda_{ij} \mathbb{E}_q[b_i + b_j] \right) \right) \mathbb{E}_q[\mathbf{v}_j] - \mathbb{E}_{\tilde{q}}[\log \tilde{q}(\mathbf{u}_i)], \end{aligned} \quad (14)$$

where tr denotes the trace operator. $\mathcal{L}[\tilde{q}(\mathbf{u}_i)]$ is maximized when $\tilde{q}(\mathbf{u}_i)$ is a Gaussian density $\mathcal{N}(\mathbf{u}_i; \boldsymbol{\mu}_i, \mathbf{P}_i^{-1})$ whose natural parameters \mathbf{P}_i and $\mathbf{m}_i \doteq \mathbf{P}_i \boldsymbol{\mu}_i$ are given by (6) and (7); they accompany $\mathbb{E}_q[\mathbf{u}_i \mathbf{u}_i^T]$ and $\mathbb{E}_q[\mathbf{u}_i]$ in the quadratic and linear terms above.

The above expression contains a sum over $j = 1, \dots, J$ and a further inner sum over $d' = 1, \dots, D'$. The scalable evaluation for \mathbf{P}_i and \mathbf{m}_i in Section 4 incorporates caches \mathbf{P}_\ominus , $\mathbf{m}_\ominus^\dagger$, and $\mathbf{m}_\ominus^\ddagger$, and only requires a sparse sum over $j \in \mathcal{G}(i)$. The simplification is obtained by using

1. $\xi_{ij} = \xi^*$ (and hence $\lambda_{ij} = \lambda^*$) for all $j \notin \mathcal{G}(i)$;
2. $c_{ij} = 0$ for all $j \notin \mathcal{G}(i)$;
3. $\mathbb{E}_q[y_{d'i}] = s_i$ for all $d' = 1, \dots, D'$;
4. $\mathbb{E}_q[z_{d'j}] = t_j$ for all $d' = 1, \dots, D'$.

B.2.2 Intermediate bounds

The bound $\mathcal{L}[\tilde{q}(\mathbf{u}_i)]$ is maximized at $\tilde{q}(\mathbf{u}_i) = \mathcal{N}(\mathbf{u}_i; \boldsymbol{\mu}_i, \mathbf{P}_i^{-1})$. With $q'(u_{ik}) \doteq \mathcal{N}(u_{ik}; \mu_{ik}, P_{i,kk}^{-1})$ being the minimizer of the Kullback-Leibler divergence

$$\prod_k q'(u_{ik}) \doteq \underset{\prod_k q(u_{ik})}{\operatorname{argmin}} \operatorname{D}_{\text{KL}} \left(\prod_k q(u_{ik}) \parallel \tilde{q}(\mathbf{u}_i) \right),$$

we now show that $\mathcal{L}[\tilde{q}(\mathbf{u}_i)]$ serves as a temporary or *intermediate* lower bound to $\log p(\mathcal{D})$:

$$\tilde{\mathcal{L}}[\tilde{q}(\mathbf{u}_i)] \geq \mathcal{L} \left[\prod_k q'(u_{ik}) \right]. \quad (15)$$

The bound in (15) follows by substituting $\mathbb{E}_{\tilde{q}}[\mathbf{u}_i \mathbf{u}_i^T] = \boldsymbol{\mu}_i \boldsymbol{\mu}_i^T + \mathbf{P}_i^{-1}$ in (14):

$$\begin{aligned} \tilde{\mathcal{L}}[\tilde{q}(\mathbf{u}_i)] &= -\frac{1}{2} \operatorname{tr} \mathbb{E}_{\tilde{q}}[\mathbf{u}_i \mathbf{u}_i^T] \mathbf{P}_i + \mathbb{E}_{\tilde{q}}[\mathbf{u}_i]^T \mathbf{P}_i \boldsymbol{\mu}_i - \mathbb{E}_{\tilde{q}}[\log \tilde{q}(\mathbf{u}_i)] \\ &= -\frac{K}{2} + \frac{1}{2} \boldsymbol{\mu}_i^T \mathbf{P}_i \boldsymbol{\mu}_i - \left(-\frac{K}{2} \log(2\pi e) + \frac{1}{2} \log |\mathbf{P}_i| \right). \end{aligned}$$

Let $\operatorname{diag}(\mathbf{P}_i)$ indicate the K -by- K matrix that contains only the diagonal of \mathbf{P}_i . As $\mathbb{E}_{q'}[\mathbf{u}_i \mathbf{u}_i^T] = \boldsymbol{\mu}_i \boldsymbol{\mu}_i^T + \operatorname{diag}(\mathbf{P}_i)^{-1}$ and $\operatorname{tr} \operatorname{diag}(\mathbf{P}_i)^{-1} \mathbf{P}_i = K$, the second bound expands as

$$\begin{aligned} \mathcal{L} \left[\prod_k q'(u_{ik}) \right] &= -\frac{1}{2} \operatorname{tr} \mathbb{E}_{q'}[\mathbf{u}_i \mathbf{u}_i^T] \mathbf{P}_i + \mathbb{E}_{q'}[\mathbf{u}_i]^T \mathbf{P}_i \boldsymbol{\mu}_i - \mathbb{E}_{q'} \left[\sum_k \log q'(u_{ik}) \right] \\ &= -\frac{K}{2} + \frac{1}{2} \boldsymbol{\mu}_i^T \mathbf{P}_i \boldsymbol{\mu}_i - \left(-\frac{K}{2} \log(2\pi e) + \frac{1}{2} \log |\operatorname{diag}(\mathbf{P}_i)| \right). \end{aligned}$$

Finally, (15) follows from the identity $|\mathbf{P}_i| \leq \prod_k P_{i,kk} = |\operatorname{diag}(\mathbf{P}_i)|$ as \mathbf{P}_i is positive definite.

B.2.3 The advantage of an intermediate bound

By first solving for $\tilde{q}(\mathbf{u}_i)$, the updates in (6) and (7) require *one* sum over $j \in \mathcal{G}(i)$, and an $\mathcal{O}(K^3)$ matrix inverse to obtain $\boldsymbol{\mu}_i \doteq \mathbf{P}_i^{-1} \mathbf{m}_i$ and $\prod_k q(u_{ik})$. On the other hand, one may solve for each $q(u_{ik})$ for $k = 1, \dots, K$ in turn. Each of these K updates require a sum over $j \in \mathcal{G}(i)$, but does not require the matrix inverse. There is therefore a computational trade-off between these two options. The trade-off depends on $|\mathcal{G}(i)|$ and K , and wasn't investigated further in the paper.

B.3 Logistic bound parameter updates

All the ξ_{ij} parameters are tied to ξ^* for $(i, j) \notin \mathcal{G}$, and we write (13) as a function of ξ^* as

$$\begin{aligned} \mathcal{L}(\xi^*) &= \log \sigma(\xi^*) \sum_{(i,j) \notin \mathcal{G}} D' s_i t_j - \lambda(\xi^*) \sum_{(i,j) \notin \mathcal{G}} D' s_i t_j \mathbb{E}_q \left[(\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j)^2 \right] \\ &\quad + \left(\lambda(\xi^*) \xi^{*2} - \frac{\xi_{ij}}{2} \right) \sum_{(i,j) \notin \mathcal{G}} D' s_i t_j. \end{aligned}$$

(Notice that for $(i, j) \notin \mathcal{G}$ we have $c_{ij} = 0$, and c_{ij} does not explicitly occur in the above expression.) Recalling that $\lambda(\xi) \doteq \frac{1}{2\xi} [\sigma(\xi) - \frac{1}{2}]$ and that $\sigma(\xi) \doteq (1 + e^{-\xi})^{-1}$, the above derivative is

$$\frac{\partial \mathcal{L}(\xi^*)}{\partial \xi^*} = -\lambda'(\xi^*) \sum_{(i,j) \notin \mathcal{G}} D' s_i t_j \mathbb{E}_q \left[(\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j)^2 - (\xi^*)^2 \right].$$

As the bound is symmetric around $\xi^* = 0$ and as $\lambda'(\xi^*)$ is a monotonic function of ξ^* for $\xi^* \geq 0$, the derivative is zero when

$$(\xi^*)^2 = \frac{1}{\sum_{(i,j) \notin \mathcal{G}} s_i t_j} \sum_{(i,j) \notin \mathcal{G}} s_i t_j \mathbb{E}_q \left[(\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j)^2 \right]. \quad (16)$$

Unfortunately (16) requires a sum over $(i, j) \notin \mathcal{G}$. However, (11) states that $\xi_{ij}^2 = \mathbb{E}_q[(\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j)^2]$ can be computed and discarded for $(i, j) \in \mathcal{G}$ if required, and hence the required sum can be written in terms of cached quantities through

$$\sum_{(i,j) \notin \mathcal{G}} s_i t_j \mathbb{E}_q \left[(\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j)^2 \right] = \sum_{i,j} s_i t_j \mathbb{E}_q \left[(\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j)^2 \right] - \sum_{(i,j) \in \mathcal{G}} s_i t_j \xi_{ij}^2,$$

and using

$$\sum_{i,j} s_i t_j \mathbb{E}_q \left[(\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j)^2 \right] = \text{tr} \mathbf{P}_\oplus \mathbf{P}_\ominus + 2\mathbf{m}_\oplus^{\dagger T} \mathbf{m}_\ominus^\dagger + 2\mathbf{m}_\oplus^{\dagger T} \mathbf{m}_\ominus^\dagger + 2\nu_\oplus \nu_\ominus + \kappa_\oplus + \kappa_\ominus.$$

B.4 Categorical updates

We want to find $q(\mathbf{y}_{d'}) = \prod_{i=1}^I s_i^{y_{d'i}}$ which is a categorical distribution parameterized by \mathbf{s} . Using the notation

$$\Omega_{ij} \doteq \log \sigma(\xi_{ij}) - \lambda(\xi_{ij}) \left(\mathbb{E}_q[(\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j)^2] - \xi_{ij}^2 \right) - \frac{\xi_{ij}}{2} - \frac{1}{2} \mathbb{E}_q[\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j]$$

from Section 4, substitute $\mathbb{E}_q[y_{d'i}] = s_i$ into (13) to obtain $\mathcal{L}_\xi[q]$ as a function of \mathbf{s} :

$$\mathcal{L}^l(\mathbf{s}) = D' \sum_i s_i \sum_j t_j \Omega_{ij} + D' \sum_i s_i \mathbb{E}_q[\log \pi_i] - D' \sum_i s_i \log s_i + l \left(\sum_i s_i - 1 \right).$$

The above function includes a Lagrange multiplier l as $\sum_i s_i$ normalizes to one. The gradient of $\mathcal{L}^l(\mathbf{s})$ with respect to s_i is zero when

$$\log s_i = \mathbb{E}_q[\log \pi_i] + \sum_j t_j \Omega_{ij} + \frac{l}{D'} - 1,$$

while the Lagrange multiplier gives the normalizer so that

$$s_i = \frac{e^{\mathbb{E}_q[\log \pi_i] + \sum_j t_j \Omega_{ij}}}{\sum_{i'=1}^I e^{\mathbb{E}_q[\log \pi_{i'}] + \sum_j t_j \Omega_{i'j}}}. \quad (17)$$

B.4.1 Using caches

Evaluating $\sum_j t_j \Omega_{ij}$ in (17) for every $i = 1, \dots, I$ again leaves us with an undesirable $\mathcal{O}(IJ)$ complexity. Here, too, we shall make heavy use of cached quantities to simplify this computation. First note that

$$\sum_{j \notin \mathcal{G}(i)} t_i \Omega_{ij} = \sum_{j \notin \mathcal{G}(i)} t_i \left(\log \sigma(\xi^*) - \lambda^* \left(\mathbb{E}_q[(\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j)^2] - (\xi^*)^2 \right) - \frac{\xi^*}{2} - \frac{1}{2} \mathbb{E}_q[\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j] \right)$$

where $\lambda^* \doteq \lambda(\xi^*)$, and that $\sum_{j \notin \mathcal{G}(i)} t_i \Omega_{ij} = \sum_{j \notin \mathcal{G}(i)} t_i \Omega_{ij}^*$. We therefore compute the full sum $\sum_j t_j \Omega_{ij}^*$ using caches, and then only loop over the sparse set $j \in \mathcal{G}(i)$ to incorporate the difference. That is,

$$\begin{aligned} \sum_j t_j \Omega_{ij}^* &= -\lambda^* \left(\text{tr} \mathbb{E}_q[\mathbf{u}_i \mathbf{u}_i^T] \mathbf{P}_\ominus + 2\mathbb{E}_q[b_i \mathbf{u}_i^T] \mathbf{m}_\ominus^\dagger + 2\mathbb{E}_q[\mathbf{u}_i^T] \mathbf{m}_\ominus^\dagger + \mathbb{E}_q[b_i^2] + 2\mathbb{E}_q[b_i] \nu_\ominus + \kappa_\ominus \right) \\ &\quad + \log \sigma(\xi^*) + \frac{(\xi^*)^2}{2} \lambda^* - \frac{\xi^*}{2} - \frac{1}{2} \left(\mathbb{E}_q[\mathbf{u}_i^T] \mathbf{m}_\ominus^\dagger + \mathbb{E}_q[b_i] + \nu_\ominus \right) \end{aligned}$$

is computed using bookkeeping, and finally

$$\sum_j t_j \Omega_{ij} = \sum_{j \in \mathcal{G}(i)} t_j (\Omega_{ij} - \Omega_{ij}^*) + \sum_j t_j \Omega_{ij}^*$$

then relies on a sparse sum.




Figure 5: The factors $q(\mathbf{z}_{d'})$ and $q(\psi)$ from a model with 2.5×10^8 Windows 8 Phone App install signals as feedback pairs. With “variational model pruning” of small components, a result that is more useful in a real system can be created by using constant messages with $\mathbb{E}[\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j] = 0$ when optimizing for $q(\mathbf{y}_{d'})$ and $q(\mathbf{z}_{d'})$.

C Practical considerations

“Variational model pruning” [11] can be observed on the fully optimized $q(\mathbf{y}_{d'})$ and $q(\mathbf{z}_{d'})$ factors. In Figure 5, one sees the average $\mathbb{E}[z_{d'j}]$ tailing roughly where $c_j < 50$. The net effect of disproportionately decreasing the expected appearance probability is that c_{ij} is explained by a much larger bias b_j .

In the context of the large-scale online system in which this model is deployed, we’ve found it beneficial to substitute a constant $\mathbb{E}[\mathbf{u}_i^T \mathbf{v}_j + b_i + b_j] = 0$ when optimizing for $q(\mathbf{y}_{d'})$ and $q(\mathbf{z}_{d'})$.

D Further evaluations

As we do not directly maximize the softmax likelihood in (1), we are additionally interested in how the model’s predictions differ from those obtained from a full softmax model. This is evaluated on a much smaller scale here.

Co-authorship networks This paper’s starting point was the bilinear softmax likelihood function in (1). We will now turn to examine how much the approximation to $p(j|i, \mathcal{D})$ in (5) deviates from that of a maximum a posteriori (MAP) solution to the softmax likelihood. As discussed earlier, the softmax MAP estimate is expensive to find, we thus use the relatively smaller NIPS 1–12 co-authorship dataset² (even though it is not naturally bipartite data). We removed all single-authors which left us with $I = 1897$ authors, and treat co-authorship as *symmetric* counts in \mathcal{D} . Biases were included in (1), and excluded from our model, so that with $K = 5$ both models have the same number of parameters. We had to add the additional constraint $\mathbf{u}_i = \mathbf{v}_i$ for all i to enforce the softmax point estimate to be symmetric. This was not required for Algorithm 1, which found a symmetric solution with and without such a constraint. Figure 6 shows the predicted co-authors for A. Smola, with the top 25 predictions labelled for each model. This is a density estimation problem with scarce data and an abundance of parameters, and with no shrinkage there are many singularities in the likelihood function. With shrinkage ($\tau_u = 1$) the smallest softmax odds are 10^{-8} in Figure 6, and the small data set is memorized by the MAP solution, which might not generalize. This result underscores the need for a Bayesian approach. Although the most probable predictions are still anecdotally interpretable, we note that a truer comparison would be against posterior predictions that are estimated using Markov chain Monte Carlo samples with (1) as likelihood, but leave this research to future work.

²www.autonlab.org/autonweb/17433

- [14] A. Mnih and R. Salakhutdinov. Probabilistic matrix factorization. In *Advances in Neural Information Processing Systems 20*, pages 1257–1264, 2008.
- [15] A. Mnih and Y. W. Teh. A fast and simple algorithm for training neural probabilistic language models. In *Proceedings of the 29th International Conference on Machine Learning*, pages 1751–1758, 2012.
- [16] A. Mnih and Y. W. Teh. Learning label trees for probabilistic modelling of implicit feedback. In *Advances in Neural Information Processing Systems 25*, pages 2825–2833, 2012.
- [17] R. Pan and M. Scholz. Mind the gaps: Weighting the unknown in large-scale one-class collaborative filtering. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 667–675, 2009.
- [18] U. Paquet and N. Koenigstein. One-class collaborative filtering with random graphs. In *Proceedings of the 22nd International Conference on World Wide Web*, pages 999–1008, 2013.
- [19] U. Paquet and N. Koenigstein. One-class collaborative filtering with random graphs: Annotated version. *CoRR*, abs/1309.6786, 2013.
- [20] U. Paquet, B. Thomson, and O. Winther. A hierarchical model for ordinal matrix factorization. *Statistics and Computing*, 22(4):945–957, 2012.
- [21] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme. BPR: Bayesian personalized ranking from implicit feedback. In *Uncertainty in Artificial Intelligence*, pages 452–461, 2009.
- [22] R. Salakhutdinov and A. Mnih. Bayesian probabilistic matrix factorization using Markov chain Monte Carlo. In *Proceedings of the 25th International Conference on Machine Learning*, pages 880–887, 2008.
- [23] V. Sindhwani, S. S. Bucak, J. Hu, and A. Mojsilovic. One-class matrix completion with low-density factorizations. In *IEEE 10th International Conference on Data Mining*, pages 1055–1060, 2010.
- [24] S. R. Waterhouse, D.J.C. MacKay, and A. J. Robinson. Bayesian methods for mixtures of experts. In *Advances in Neural Information Processing Systems 8*, pages 351–357, 1996.